

Simplify Your Systems with
Scale Computing HyperCore™

SC//HYPERCORE
PLATFORM

Move Beyond Traditional IT Silos

Scale Computing HyperCore™ saves you time and valuable resources because your software, servers, and storage are in a fully integrated platform. The same innovative software and simple user interface power your infrastructure regardless of your hardware configuration.

Using patented HyperCore™ technology, the award-winning self-healing platform identifies, reduces, and corrects problems in real-time. Achieve results easier and faster, even when local IT resources are scarce. SC//HyperCore makes ensuring application uptime easier for IT to manage and for customers to afford.

SIMPLE


Deploy and manage easily. With the self-healing capabilities and intelligent automation of SC//HyperCore, customers spend less time on infrastructure maintenance and focus on strategic projects to grow their business.

EFFICIENT


Eliminate the need to combine traditional virtualization software, disaster recovery software, servers, and shared storage from separate vendors to create a virtualized environment. SC//HyperCore's lightweight, all-in-one architecture makes it easy to deploy fully integrated, highly available virtualization right out of the box.

FUTURE PROOF

Plan the perfect future by not predicting it. Capacity, performance requirements, and hardware specs will forever change. Legacy applications will remain mission-critical. Modern applications require new architectures. Simply mix and match the old and the new on SC//HyperCore for a future-proof environment that can scale up or down as needed.



SC//HyperCore heads-up display



KEY BENEFITS

TAKE IT EMERGENCIES OUT OF THE EQUATION

Scale Computing's self-healing technology automatically corrects issues, so you can avoid an IT crisis from occurring at just the wrong time. Keep systems up to date, and repair failures as part of a regularly scheduled maintenance cycle.

REDUCE YOUR FOOTPRINT

Use fewer resources, while cutting major costs. The overall efficiency of SC//HyperCore allows the use of much smaller platforms, while still reserving resources to run user workloads. This provides significant savings when multiplied across hundreds or thousands of edge sites.

NO NEED FOR TRADITIONAL STORAGE MANAGEMENT

Secure your data. Reduce latency while eliminating the need for traditional storage management. The Scale Computing Reliable Independent Block Engine (SCRIBE) combines SC//HyperCore storage drives into a single storage pool, requiring zero user administration, and is available to all nodes of the system without requiring any file systems, protocols, or VSAs.

HIGHLIGHTS

Software-Defined Storage

All components—storage, virtualization, software and hardware—interface directly through the SC//HyperCore hypervisor and SCRIBE storage layers to create an ideal computing platform that can be deployed anywhere from the data center to the edge of the network.

- Enable configurable SSD priority allocation at the individual virtual disk-level and intelligent data block priority based on block I/O heat mapping
- Discover all block storage devices— including flash-based solid-state disks (SSDs) and conventional spinning disks (SATA or SAS)
- Aggregate block storage devices across all nodes of SC//HyperCore into a single managed pool of storage
- Allow sophisticated data redundancy, load balancing intelligence, and I/O tiered prioritization
- Efficiently use flash storage when available for tiered data placement

Software-Managed Compute

SC//HyperCore is a lightweight, type 1 (bare metal) hypervisor that is directly integrated into the OS kernel and leverages the virtualization offload capabilities provided by modern CPU architectures. Specifically, SC//HyperCore is based on components of the KVM hypervisor, which has been part of the Linux mainline kernel for many years and has been extensively field-proven in large-scale environments.

- Run VMs on SC//HyperCore to have direct block-level access to the SCRIBE virtual storage device (VSD) virtual disks in the clustered storage pool without the complexity or performance overhead introduced by using remote storage protocols and accessing remote storage over a network
- Allow direct data flow to benefit from zero-copy shared memory performance—unlike other architectures the storage layer runs parallel to the hypervisor, instead of inside a VM as a VSA or controller VM

“Scale Computing HyperCore is unquestionably one of the top solutions for saving time and providing several job benefits. It is much easier to set up, use, and clone templates, and no one needs to go through extensive training to learn how to use it, nor do we need an expert to assist us to manage our system because this program allows our employees to handle things on their own. Setting up installation is now quite straightforward, making it much easier for our group to distribute new goods in a timely way.”



Mix and Match Nodes

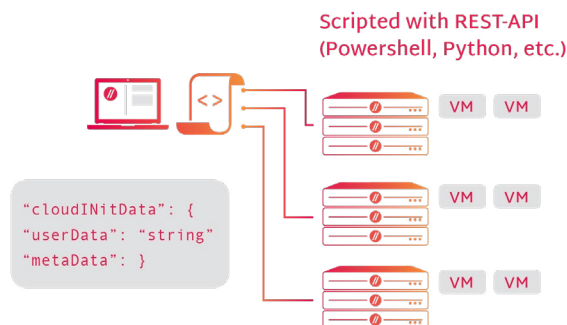
Hardware appliances of different sizes can be added to and coexist in the same cluster and contribute to the cluster storage pool. Even appliances with dissimilar storage such as all HDD, hybrid HDD + SSD, or all SSDs can be combined as a cluster grows. This allows clusters to grow not just linearly, but in whatever way is needed or desired for application workloads.

- Facilitate non-disruptive upgrades to the infrastructure by adding newer, bigger nodes and then retiring or repurposing older small nodes if desired, all while keeping the VMs running
- Initialize a cluster with just a single node, and scale up to eight nodes
- Handle your smallest edge and largest data center workloads in one unified manner

REST-APIs Complete with Cloud-Init

SC//HyperCore includes REST-APIs that enhance the speed and ease with which users can deploy virtual machines (VMs) at scale using cloud-init.

- Use common VM templates and provide them with their unique configuration information at first boot via scripting
- Allow easy, mass provisioning of customized VMs
- Avoid the need to manually create and individually customize VMs, but programmatically provide hundreds or thousands of machines with their own settings via script



Containers Support

To run containers on SC//HyperCore, simply deploy a container-optimized operating system with a container runtime of choice (often Docker, or in a Kubernetes environment, containerd or CRI-O).

- REST-APIs and support of cloud-init radically improve how users run containerized workloads
- Automated installation of the operating system, container runtime, and workload containers
- Enable consistent change control and more reliable update through standardization

